

Sub
A1

Hardware for Use with Compiler Generated Branch Information

BACKGROUND

5 1. Field of the Present Invention

The present invention generally relates to the field of microprocessors and more particularly to a microprocessor including hardware designed to minimize branch misprediction by using compiler generated branch information to speculatively execute instructions following a
10 branch condition.

2. History of Related Art

A major challenge for designers of gigahertz microprocessors is to take advantage of state-of-the-art technologies while maintaining compatibility with the enormous base of installed
15 software designed for operation with a particular instruction set architecture (ISA). To address this problem, designers have implemented "layered architecture" microprocessors that are adapted to receive instructions formatted according to an existing ISA and to convert the instruction format of the received instructions to an internal ISA that is more suitable for operation in gigahertz execution units.

20 Because a layered architecture adds to the processor pipeline and increases that number of instructions that are potentially "in flight" at any given time, the branch mispredict penalty associated with a layered architecture is of great concern. Dynamic mechanisms for predicting branches can be highly accurate in predicting certain types of branches, such as the branches associated with a loop that is executed multiple times. For a frequently encountered class of
25 branches (referred to herein as random branches) such as a branch predicated upon a comparison between the contents of a random memory location and the contents of a register, however, conventional prediction mechanisms are not typically effective in reducing the misprediction rate significantly below 50%. Because of the large misprediction penalty associated with superscalar

architectures in general and layered architecture processors in particular, it would be highly desirable to address the expensive misprediction penalty associated with random branches.

5

SUMMARY OF THE INVENTION

The problem identified above is addressed by a method of executing microprocessor instructions and an associated microprocessor according to the present invention. Initially, a conditional branch instruction is fetched from a storage unit such as an instruction cache. A
10 fetch unit of the microprocessor detects branch prediction information embedded in the branch instruction. Depending upon the state of the branch prediction information, instructions from the branch-taken path and the branch-not-taken path of the branch instruction are fetched. The branch-not-taken path instructions and the branch-taken path instruction may be speculatively executed. Upon executing the conditional branch instruction, the speculative results from the branch-taken path are discarded if the branch is not taken and speculative results from the branch-not-taken path are discarded if the branch is taken. The branch prediction information may include compiler-generated information indicative of the context in which the conditional branch instruction is used. In one embodiment, the branch prediction information causes instruction fetching from both the taken and not taken branches if the compiler determines the
15 branch instruction to unpredictable. In another embodiment, fetching instructions from the branch-taken path includes fetching a predetermined number of instructions from the branch-taken path and a predetermined number of instructions from the branch-not-taken path. In another embodiment, instructions are fetch down the branch-not-taken path until a subsequent branch instruction is encountered.

20
25

BRIEF DESCRIPTION OF THE DRAWINGS

Sub A3
Other objects and advantages of the invention will become apparent upon reading the following detailed description and upon reference to the accompanying drawings in which:

A3
FIG 1 is a block diagram of a data processing system;

FIG 2 is a block diagram illustrating selected features of a processor suitable for use in the data processing system of FIG 1;

FIG 3 is an exemplary code segment suitable for implementing the conditional branch
5 information of the present invention;

FIG 4 illustrates a branch conditional statement according to the present invention; and

FIG 5 is a block diagram illustrating features of a branch prediction unit according to one embodiment of the present invention.

While the invention is susceptible to various modifications and alternative forms, specific
10 embodiments thereof are shown by way of example in the drawings and will herein be described in detail. It should be understood, however, that the drawings and detailed description presented herein are not intended to limit the invention to the particular embodiment disclosed, but on the contrary, the intention is to cover all modifications, equivalents, and alternatives falling within the spirit and scope of the present invention as defined by the appended claims.

DETAILED DESCRIPTION OF A PREFERRED EMBODIMENT
OF THE PRESENT INVENTION

Referring now to FIG 1, an embodiment of a data processing system 100 according to the
20 present invention is depicted. System 100 has one or more central processing units (processors) 101a, 101b, 101c, etc. (collectively or generically referred to as processor(s) 101. In one embodiment, each processor 101 may comprise a reduced instruction set computer (RISC) microprocessor. Additional information concerning RISC processors in general is available in C. May et al. Ed., *PowerPC Architecture: A Specification for a New Family of RISC Processors*,
25 (Morgan Kaufmann, 1994 2d edition). Processors 101 are coupled to system memory 250 and various other components via system bus 113. Read only memory (ROM) 102 is coupled to the system bus 113 and may include a basic input/output system (BIOS), which controls certain basic functions of system 100. FIG 1 further depicts an I/O adapter 107 and a network adapter 106 coupled to the system bus 113. I/O adapter 107 may be a small computer system interface

(SCSI) adapter that communicates with a hard disk 103 and/or tape storage drive 105. I/O adapter 107, hard disk 103, and tape storage device 105 are collectively referred to herein as mass storage 104. A network adapter 106 interconnects bus 113 with an outside network enabling data processing system 100 to communicate with other such systems. Display monitor 136 is connected to system bus 113 by display adapter 112, which may include a graphics adapter to improve the performance of graphics intensive applications and a video controller. In one embodiment, adapters 107, 106, and 112 may be connected to one or more I/O busses that are connected to system bus 113 via an intermediate bus bridge (not shown). Suitable I/O busses for connecting peripheral devices such as hard disk controllers, network adapters, and graphics adapters include the Peripheral Components Interface (PCI) bus according to PCI Local Bus Specification Rev. 2.2 available from the PCI Special Interest Group, Hillsboro OR, and incorporated by reference herein. Additional input/output devices are shown as connected to system bus 113 via user interface adapter 108 and display adapter 112. A keyboard 109, mouse 110, and speaker 111 all interconnected to bus 113 via user interface adapter 108, which may include, for example, a SuperI/O chip integrating multiple device adapters into a single integrated circuit. For additional information concerning one such chip, the reader is referred to the PC87338/PC97338 ACPI 1.0 and PC98/99 Compliant SuperI/O data sheet from National Semiconductor Corporation (November 1998) at www.national.com. Thus, as configured in FIG 1, system 100 includes processing means in the form of processors 101, storage means including system memory 250 and mass storage 104, input means such as keyboard 109 and mouse 110, and output means including speaker 111 and display 136. In one embodiment a portion of system memory 250 and mass storage 104 collectively store an operating system such as the AIX® operating system from IBM Corporation to coordinate the functions of the various components shown in FIG 1. Additional detail concerning the AIX operating system is available in *AIX Version 4.3 Technical Reference: Base Operating System and Extensions, Volumes 1 and 2* (order numbers SC23-4159 and SC23-4160); *AIX Version 4.3 System User's Guide: Communications and Networks* (order number SC23-4122); and *AIX Version 4.3 System User's Guide: Operating System and Devices* (order number SC23-4121) from IBM Corporation at www.ibm.com and incorporated by reference herein.

Turning now to FIG 2, a simplified block diagram of a processor **101** according to one embodiment of the present invention is illustrated. Processor **101** as depicted in FIG 2 includes an instruction fetch unit **202** suitable for generating an address of the next instruction to be fetched. The fetched instruction address generated by fetch unit **202** is loaded into a next instruction address latch **204** and provided to an instruction cache **210**.

Fetch unit **202** further includes branch prediction logic **206**. As its name suggests, branch prediction logic **206** is adapted to predict the outcome of a decision that effects the program execution flow. The ability to correctly predict branch decisions is a significant factor in the overall ability of processor **101** to achieve improved performance by executing instructions speculatively and out-of-order.

The address produced by fetch unit **202** is provided to an instruction cache **210**, which contains a subset of the contents of system memory in a high-speed storage facility. If the address instruction generated by fetch unit **202** corresponds to a system memory location that is currently replicated in instruction cache **210**, instruction cache **210** forwards the corresponding instruction to cracking logic **212**. If the instruction corresponding to the instruction address generated by fetch unit **202** does not currently reside in instruction cache **210**, the contents of instruction cache **210** must be updated with the contents of the appropriate locations in system memory before the instruction can be forwarded to cracking logic **212**. Additional details of instruction fetch unit **202** are described in greater detail below with respect to FIG 3.

Cracking logic **212** is adapted to modify an incoming instruction stream to produce a set of instructions optimized for executing in an underlying execution unit at extremely high operating frequencies (i.e., operating frequencies exceeding 1 GHz). Cracking logic **212** may, for example, receive instructions in a 32-bit wide format such as instructions supported by the PowerPC® instruction set. Detailed information regarding the PowerPC® instruction set is available in the *PowerPC® Microprocessor Family: The Programming Environments for 32-Bit*

Microprocessors available from IBM Corporation. (Order No. G522-0290-01), which is incorporated by reference herein.

The format of the instructions generated by cracking logic 212 may include explicit fields for information that is merely implied in the format of the fetched instructions such that the format of instructions generated by cracking logic 212 is wider than the format of instructions. In one embodiment, for example, the fetched instructions are encoded according to a 32-bit instruction format and the format of instructions generated by cracking logic 212 is 64 or more bits wide. Cracking logic 212 is designed to generate these wide instructions according to a predefined set of cracking rules. In addition, cracking unit 212 may logically organize a set of instructions into one or more instruction groups to simply exception and completion handling logic. Instruction groups are disclosed in greater detail in U.S. Patent Application No. 09/428,399 entitled *Instruction Group Organization and Exception Handling in a Microprocessor*, filed October 28, 1999, which shares a common assignee with the present application and is incorporated by reference herein.

Returning now to FIG 2, the wide instructions generated in the preferred embodiment of cracking unit 212 are forwarded to dispatch unit 214. Dispatch unit 214 is responsible for determining which instructions are capable of being executed and forwarding these executable instructions to issue queues 220. In addition, dispatch unit 214 communicates with dispatch and completion control logic 216 to keep track of the order in which instructions were issued and the completion status of these instructions to facilitate out-of-order execution. In an embodiment of processor 101 in which cracking unit 212 organizes incoming instructions into instruction groups, each instruction group is assigned a group tag (GTAG) by completion and control logic 216 that conveys the ordering of the issued instruction groups. As an example, dispatch unit 214 may assign monotonically increasing values to consecutive instruction groups. With this arrangement, instruction groups with lower GTAG values are known to have issued prior to (i.e., are older than) instruction groups with larger GTAG values. In association with dispatch and completion control logic 216, a completion table 218 is utilized in one embodiment of the

present invention to track the status of issued instruction groups.

In the embodiment of processor 101 depicted in FIG 2, instructions are issued from dispatch unit 214 to issue queues 220 where they await execution in corresponding execution units 222. Processor 101 may include a variety of types of executions pipes, each designed to execute a subset of the processor's instruction set. In one embodiment, execution units 222 may include a branch unit 224, a load store unit 226, a fixed-point arithmetic unit 228, and a floating point unit 230. Each execution unit 222 may comprise two or more pipeline stages.

Instructions stored in issue queues 220 may be issued to execution units 222 using any of a variety of issue priority algorithms. In one embodiment, for example, the oldest pending instruction in an issue queue 220 is the next instruction issued to execution units 222. In this embodiment, the GTAG values assigned by dispatch unit 214 are utilized to determine the relative age of instructions pending in the issue queues 220.

Prior to issue, the destination register operand of each instruction is assigned to an available rename register. When an instruction is ultimately forwarded from issue queues 120 to the appropriate execution unit, the execution unit performs the operation indicated by the instruction's opcode and writes the instruction's result to the instruction's rename register by the time the instruction reaches a finish stage (indicated by reference numeral 132) of the pipeline. A mapping is maintained between the rename registers and their corresponding architected registers. When all instructions in an instruction group (and all instructions in older instruction groups) finish without generating an exception, a completion pointer in the completion table 218 is incremented to the next instruction group.

When the completion pointer is incremented to a new instruction group, the rename registers associated with the instructions in the old instruction group are released thereby committing the results of the instructions in the old instruction group. If one or more instructions older than a finished (but not yet committed) instruction generate an exception, the instruction

generating the exception and all younger instructions are flushed and a rename recovery routine is invoked to return the rename register mapping to the last known valid state. In this manner, processor 101 enables speculative and out-of-order instruction execution.

5 Speculative instruction execution is most frequently encountered in conjunction with conditional branches. When instruction fetch unit 202 fetches a conditional branch instruction from instruction cache 210, the fetch unit 202 cannot know with certainty which instruction will be executed immediately after the conditional branch is executed because the condition on which the branch is dependent is evaluated when the conditional branch is executed. To prevent an
10 instruction fetch stall from occurring following every conditional branch instruction, branch prediction unit 206 is included in fetch unit 202 to predict the "outcome" of conditional branches.

34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65
66
67
68
69
70
71
72
73
74
75
76
77
78
79
80
81
82
83
84
85
86
87
88
89
90
91
92
93
94
95
96
97
98
99
100
101
102
103
104
105
106
107
108
109
110
111
112
113
114
115
116
117
118
119
120
121
122
123
124
125
126
127
128
129
130
131
132
133
134
135
136
137
138
139
140
141
142
143
144
145
146
147
148
149
150
151
152
153
154
155
156
157
158
159
160
161
162
163
164
165
166
167
168
169
170
171
172
173
174
175
176
177
178
179
180
181
182
183
184
185
186
187
188
189
190
191
192
193
194
195
196
197
198
199
200
201
202
203
204
205
206
207
208
209
210
211
212
213
214
215
216
217
218
219
220
221
222
223
224
225
226
227
228
229
230
231
232
233
234
235
236
237
238
239
240
241
242
243
244
245
246
247
248
249
250
251
252
253
254
255
256
257
258
259
260
261
262
263
264
265
266
267
268
269
270
271
272
273
274
275
276
277
278
279
280
281
282
283
284
285
286
287
288
289
290
291
292
293
294
295
296
297
298
299
300
301
302
303
304
305
306
307
308
309
310
311
312
313
314
315
316
317
318
319
320
321
322
323
324
325
326
327
328
329
330
331
332
333
334
335
336
337
338
339
340
341
342
343
344
345
346
347
348
349
350
351
352
353
354
355
356
357
358
359
360
361
362
363
364
365
366
367
368
369
370
371
372
373
374
375
376
377
378
379
380
381
382
383
384
385
386
387
388
389
390
391
392
393
394
395
396
397
398
399
400
401
402
403
404
405
406
407
408
409
410
411
412
413
414
415
416
417
418
419
420
421
422
423
424
425
426
427
428
429
430
431
432
433
434
435
436
437
438
439
440
441
442
443
444
445
446
447
448
449
450
451
452
453
454
455
456
457
458
459
460
461
462
463
464
465
466
467
468
469
470
471
472
473
474
475
476
477
478
479
480
481
482
483
484
485
486
487
488
489
490
491
492
493
494
495
496
497
498
499
500
501
502
503
504
505
506
507
508
509
510
511
512
513
514
515
516
517
518
519
520
521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539
540
541
542
543
544
545
546
547
548
549
550
551
552
553
554
555
556
557
558
559
560
561
562
563
564
565
566
567
568
569
570
571
572
573
574
575
576
577
578
579
580
581
582
583
584
585
586
587
588
589
590
591
592
593
594
595
596
597
598
599
600
601
602
603
604
605
606
607
608
609
610
611
612
613
614
615
616
617
618
619
620
621
622
623
624
625
626
627
628
629
630
631
632
633
634
635
636
637
638
639
640
641
642
643
644
645
646
647
648
649
650
651
652
653
654
655
656
657
658
659
660
661
662
663
664
665
666
667
668
669
670
671
672
673
674
675
676
677
678
679
680
681
682
683
684
685
686
687
688
689
690
691
692
693
694
695
696
697
698
699
700
701
702
703
704
705
706
707
708
709
710
711
712
713
714
715
716
717
718
719
720
721
722
723
724
725
726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755
756
757
758
759
760
761
762
763
764
765
766
767
768
769
770
771
772
773
774
775
776
777
778
779
780
781
782
783
784
785
786
787
788
789
790
791
792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809
810
811
812
813
814
815
816
817
818
819
820
821
822
823
824
825
826
827
828
829
830
831
832
833
834
835
836
837
838
839
840
841
842
843
844
845
846
847
848
849
850
851
852
853
854
855
856
857
858
859
860
861
862
863
864
865
866
867
868
869
870
871
872
873
874
875
876
877
878
879
880
881
882
883
884
885
886
887
888
889
890
891
892
893
894
895
896
897
898
899
900
901
902
903
904
905
906
907
908
909
910
911
912
913
914
915
916
917
918
919
920
921
922
923
924
925
926
927
928
929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971
972
973
974
975
976
977
978
979
980
981
982
983
984
985
986
987
988
989
990
991
992
993
994
995
996
997
998
999
1000
1001
1002
1003
1004
1005
1006
1007
1008
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1020
1021
1022
1023
1024
1025
1026
1027
1028
1029
1030
1031
1032
1033
1034
1035
1036
1037
1038
1039
1040
1041
1042
1043
1044
1045
1046
1047
1048
1049
1050
1051
1052
1053
1054
1055
1056
1057
1058
1059
1060
1061
1062
1063
1064
1065
1066
1067
1068
1069
1070
1071
1072
1073
1074
1075
1076
1077
1078
1079
1080
1081
1082
1083
1084
1085
1086
1087
1088
1089
1090
1091
1092
1093
1094
1095
1096
1097
1098
1099
1100
1101
1102
1103
1104
1105
1106
1107
1108
1109
1110
1111
1112
1113
1114
1115
1116
1117
1118
1119
1120
1121
1122
1123
1124
1125
1126
1127
1128
1129
1130
1131
1132
1133
1134
1135
1136
1137
1138
1139
1140
1141
1142
1143
1144
1145
1146
1147
1148
1149
1150
1151
1152
1153
1154
1155
1156
1157
1158
1159
1160
1161
1162
1163
1164
1165
1166
1167
1168
1169
1170
1171
1172
1173
1174
1175
1176
1177
1178
1179
1180
1181
1182
1183
1184
1185
1186
1187
1188
1189
1190
1191
1192
1193
1194
1195
1196
1197
1198
1199
1200
1201
1202
1203
1204
1205
1206
1207
1208
1209
1210
1211
1212
1213
1214
1215
1216
1217
1218
1219
1220
1221
1222
1223
1224
1225
1226
1227
1228
1229
1230
1231
1232
1233
1234
1235
1236
1237
1238
1239
1240
1241
1242
1243
1244
1245
1246
1247
1248
1249
1250
1251
1252
1253
1254
1255
1256
1257
1258
1259
1260
1261
1262
1263
1264
1265
1266
1267
1268
1269
1270
1271
1272
1273
1274
1275
1276
1277
1278
1279
1280
1281
1282
1283
1284
1285
1286
1287
1288
1289
1290
1291
1292
1293
1294
1295
1296
1297
1298
1299
1300
1301
1302
1303
1304
1305
1306
1307
1308
1309
1310
1311
1312
1313
1314
1315
1316
1317
1318
1319
1320
1321
1322
1323
1324
1325
1326
1327
1328
1329
1330
1331
1332
1333
1334
1335
1336
1337
1338
1339
1340
1341
1342
1343
1344
1345
1346
1347
1348
1349
1350
1351
1352
1353
1354
1355
1356
1357
1358
1359
1360
1361
1362
1363
1364
1365
1366
1367
1368
1369
1370
1371
1372
1373
1374
1375
1376
1377
1378
1379
1380
1381
1382
1383
1384
1385
1386
1387
1388
1389
1390
1391
1392
1393
1394
1395
1396
1397
1398
1399
1400
1401
1402
1403
1404
1405
1406
1407
1408
1409
1410
1411
1412
1413
1414
1415
1416
1417
1418
1419
1420
1421
1422
1423
1424
1425
1426
1427
1428
1429
1430
1431
1432
1433
1434
1435
1436
1437
1438
1439
1440
1441
1442
1443
1444
1445
1446
1447
1448
1449
1450
1451
1452
1453
1454
1455
1456
1457
1458
1459
1460
1461
1462
1463
1464
1465
1466
1467
1468
1469
1470
1471
1472
1473
1474
1475
1476
1477
1478
1479
1480
1481
1482
1483
1484
1485
1486
1487
1488
1489
1490
1491
1492
1493
1494
1495
1496
1497
1498
1499
1500
1501
1502
1503
1504
1505
1506
1507
1508
1509
1510
1511
1512
1513
1514
1515
1516
1517
1518
1519
1520
1521
1522
1523
1524
1525
1526
1527
1528
1529
1530
1531
1532
1533
1534
1535
1536
1537
1538
1539
1540
1541
1542
1543
1544
1545
1546
1547
1548
1549
1550
1551
1552
1553
1554
1555
1556
1557
1558
1559
1560
1561
1562
1563
1564
1565
1566
1567
1568
1569
1570
1571
1572
1573
1574
1575
1576
1577
1578
1579
1580
1581
1582
1583
1584
1585
1586
1587
1588
1589
1590
1591
1592
1593
1594
1595
1596
1597
1598
1599
1600
1601
1602
1603
1604
1605
1606
1607
1608
1609
1610
1611
1612
1613
1614
1615
1616
1617
1618
1619
1620
1621
1622
1623
1624
1625
1626
1627
1628
1629
1630
1631
1632
1633
1634
1635
1636
1637
1638
1639
1640
1641
1642
1643
1644
1645
1646
1647
1648
1649
1650
1651
1652
1653
1654
1655
1656
1657
1658
1659
1660
1661
1662
1663
1664
1665
1666
1667
1668
1669
1670
1671
1672
1673
1674
1675
1676
1677
1678
1679
1680
1681
1682
1683
1684
1685
1686
1687
1688
1689
1690
1691
1692
1693
1694
1695
1696
1697
1698
1699
1700
1701
1702
1703
1704
1705
1706
1707
1708
1709
1710
1711
1712
1713
1714
1715
1716
1717
1718
1719
1720
1721
1722
1723
1724
1725
1726
1727
1728
1729
1730
1731
1732
1733
1734
1735
1736
1737
1738
1739
1740
1741
1742
1743
1744
1745
1746
1747
1748
1749
1750
1751
1752
1753
1754
1755
1756
1757
1758
1759
1760
1761
1762
1763
1764
1765
1766
1767
1768
1769
1770
1771
1772
1773
1774
1775
1776
1777
1778
1779
1780
1781
1782
1783
1784
1785
1786
1787
1788
1789
1790
1791
1792
1793
1794
1795
1796
1797
1798
1799
1800
1801
1802
1803
1804
1805
1806
1807
1808
1809
1810
1811
1812
1813
1814
1815
1816
1817
1818
1819
1820
1821
1822
1823
1824
1825
1826
1827
1828
1829
1830
1831
1832
1833
1834
1835
1836
1837
1838
1839
1840
1841
1842
1843
1844
1845
1846
1847
1848
1849
1850
1851
1852
1853
1854
1855
1856
1857
1858
1859
1860
1861
1862
1863
1864
1865
1866
1867
1868
1869
1870
1871
1872
1873
1874
1875
1876
1877
1878
1879
1880
1881
1882
1883
1884
1885
1886
1887
1888
1889
1890
1891
1892
1893
1894
1895
1896
1897
1898
1899
1900
1901
1902
1903
1904
1905
1906
1907
1908
1909
1910
1911
1912
1913
1914
1915
1916
1917
1918
1919
1920
1921
1922
1923
1924
1925
1926
1927
1928
1929
1930
1931
1932
1933
1934
1935
1936
1937
1938
1939
1940
1941
1942
1943
1944
1945
1946
1947
1948
1949
1950
1951
1952
1953
1954
1955
1956
1957
1958
1959
1960
1961
1962
1963
1964
1965
1966
1967
1968
1969
1970
1971
1972
1973
1974
1975
1976
1977
1978
1979
1980
1981
1982
1983
1984
1985
1986
1987
1988
1989
1990
1991
1992
1993
1994
1995
1996
1997
1998
1999
2000
2001
2002
2003
2004
2005
2006
2007
2008
2009
2010
2011
2012
2013
2014
2015
2016
2017
2018
2019
2020
2021
2022
2023
2024
2025
2026
2027
2028
2029
2030
2031
2032
2033
2034
2035
2036
2037
2038
2039
2040
2041
2042
2043
2044
2045
2046
2047
2048
2049
2050
2051
2052
2053
2054
2055
2056
2057
2058
2059
2060
2061
2062
2063
2064
2065
2066
2067
2068
2069
2070
2071
2072
2073
2074
2075
2076
2077
2078
2079
2080
2081
2082
2083
2084
2085
2086
2087
2088
2089
2090
2091
2092
2093
2094
2095
2096
2097
2098
2099
2100
2101
2102
2103
2104
2105
2106
2107
2108
2109
2110
2111
2112
2113
2114
2115
2116
2117
2118
2119
2120
2121
2122
2123
2124
2125
2126
2127
2128
2129
2130
2131
2132
2133
2134
2135
2136
2137
2138
2139
2140
2141
2142
2143
2144
2145
2146
2147
2148
2149
2150
2151
2152
2153
2154
2155
2156
2157
2158
2159
2160
2161
2162
2163
2164
2165
2166
2167
2168
2169
2170
2171
2172
2173
2174
2175
2176
2177
2178
2179
2180
2181
2182
2183
2184
2185
2186
2187
2188
2189
2190
2191
2192
2193
2194
2195
2196
2197
2198
2199
2200
2201
2202
2203
2204
2205
2206
2207

of this branch consistently. Thus, the conditional branch instruction at IA2 is said to be unpredictable. In other words, the likelihood of correctly predicting the conditional branch is approximately 50% and the ability to improve the percentage significantly using branch history information is limited. In one embodiment of the invention, an instruction is referred to as
5 unpredictable if the compiler determines the probability of correctly predicting the branch to be less than 75%.

Under these circumstances, the conditional branch instruction carries a significant misprediction cost because the instruction will incur a branch mispredict penalty approximately
10 half the time it is executed. Branch misprediction causes microprocessor 101 to flush a potentially large number of in-flight instructions and to restore the state of the instruction address latch and the architected registers to the state that existed prior to execution of the mispredicted branch. Microprocessor 101 according to the present invention addresses the branch mispredict penalty associated with random branches (i.e., branch instructions that are not susceptible to highly accurate branch prediction), by speculatively fetching and executing both sides (i.e., the branch-taken side and the branch not take side) of a branch sequence based upon branch prediction information that is encoded in the conditional branch instruction. When the branch instruction is executed, the results from the correctly predicted side of the branch can be committed to register files while the results from the wrongly predicted side are discarded. In
15 this manner, the relatively expensive branch mispredict penalty is avoided for the relatively modest price associated with speculatively executing a few additional instructions.

Turning now to FIG 5, additional detail of branch prediction unit 206 according to one embodiment of the invention is presented. In the depicted embodiment, branch prediction unit
20 206 includes prediction logic 502 that is configured to predict branches when enabled. Branch prediction unit 206 also includes prediction bypass unit 504 that is used when processor 101 determines that branch prediction is not effective with respect to a particular conditional branch.

In the depicted embodiment, prediction logic 502 and prediction bypass unit 504 both

utilized prediction information (identified as XY in the figure) that is supplied by the branch instruction retrieved from instruction cache 210. In one embodiment, the prediction information represents compiler-generated bits of information that are incorporated into the opcode bits of the conditional branch instruction. Branch prediction unit 206 receives the branch prediction
5 information when the instruction is retrieved from the instruction cache and forwards the information to prediction logic 502 and prediction bypass unit 504. Depending on the state of the prediction information, branch prediction unit 206 uses either prediction logic 502 to predict the branch (perhaps in conjunction with branch history table 503) and forwards the prediction result (i.e., the instruction address of the predicted next instruction) back to instruction address
10 latch 204 or uses prediction bypass unit 504 to issue instruction addresses representing both sides of the branch under consideration.

Referring momentarily to FIG 4, a conditional branch instruction formatted in accordance with the PowerPC® instruction set and suitable for use with the present invention is presented. In the illustrated example, branch conditional (BC) instruction 400 includes a 6-bit primary opcode field 402, and a 5-bit secondary opcode field 404. Secondary opcode field 404 indicates whether the branch is taken if the appropriate condition (represented by a bit in the condition register) is true or false. In addition, the depicted embodiment of secondary opcode field 404 includes a pair of branch prediction information bits, identified as the XY bits. The XY bits are
20 generated during compilation of the executable code based upon the compiler's interpretation of the context in which the conditional branch statement is used. If the compiler determines that branch prediction is unlikely to significantly improve the branch prediction rate, the compiler can encode the XY bits with an appropriate value to indicate that branch prediction is to be bypassed. Similarly, if the compiler determines from its context that a particular branch conditional
25 instruction is susceptible to accurate branch prediction, the compiler can encode the XY bits of the branch instruction accordingly.

Returning now to FIG 5, when the instruction address of a conditional branch instruction is clocked out of instruction address latch 204, it is presented to instruction cache 210. In

response, instruction cache 210 provides the appropriate instruction to the execution dispatch / execution pipeline below (not depicted in FIG 5). In addition, the branch prediction information bits (the XY bits) are provided to prediction logic 502 and prediction bypass unit 504 of branch prediction unit 206. Appropriate latching circuitry may be included in branch prediction unit 206
5 to ensure that the conditional branch's instruction address arrives at prediction logic 502 and prediction bypass unit 504 at the same time as the instruction's branch prediction information.

If the branch prediction information indicates that branch prediction is likely to be ineffective, prediction bypass unit 504 is used to cause instructions on both sides of the
10 conditional branch to be fetched and provided to the execution pipeline. An example of the operation of branch prediction unit 206 is illustrated in FIG 5 with respect to the code segment presented in FIG 3. When the instruction address (IA2) of the conditional branch instruction comes out of instruction latch 204, it retrieves the 32-bit instruction from instruction cache 210 and forwards the instruction to the underlying dispatch / execution circuitry. In addition, the
15 instruction's branch prediction bits are forwarded to prediction bypass unit 504 and prediction logic 502. Because the particular branch conditional instruction illustrated in FIG 3 is used in a context in which branch prediction is not likely to be effective, the compiler has presumably encoded the branch prediction bits of the branch conditional instruction with the appropriate code to invoke branch prediction bypass unit 504.

Upon receiving the XY bits from instruction cache 210 (at the same time as it receives the branch conditional instruction address), branch prediction bypass unit 504 generates a sequence of instruction addresses representing both sides of the branch conditional sequence. To facilitate the generation of instruction addresses, prediction bypass unit 504 may receive branch address
25 information from instruction cache 210 and may include an adder circuit for calculating the branch-taken address. Thus, branch prediction bypass unit 504 outputs a sequence of instruction addresses that causes fetch unit 202 to fetch instructions from the branch-not-taken leg of the branch conditional segment (IA3, IA4, IA5, and IA6) and from the branch-taken leg (IA7, IA8, IA9, and IA10). In one embodiment, branch prediction bypass unit 504 fetches a predetermined

number of instructions from the branch-taken path of the branch instruction and a predetermined number of instructions from the branch-not-taken path upon detecting an unpredictable branch. In another embodiment, instructions are fetched down the branch-not-taken path until a subsequent branch instruction is encountered. Thus, in the illustrated example, one embodiment of branch prediction bypass unit 504 might simply fetch three (or some other predetermined number) of instructions down both sides of the branch. In another embodiment, bypass unit 504 may fetch instructions down the branch-not-taken path (i.e., the path beginning at IA3) until the subsequent branch instruction at IA6 is encountered. In this manner, compiler generated branch prediction information is used to minimize the branch mispredict penalty in highly pipelined microprocessors thereby potentially improving processor performance.

It will be apparent to those skilled in the art having the benefit of this disclosure that the present invention contemplates improved microprocessor performance by incorporating prediction information into conditional branch instruction and using the prediction information to speculatively execute both sides of a branch when prediction is unlikely to yield consistently accurate results. It is understood that the form of the invention shown and described in the detailed description and the drawings are to be taken merely as presently preferred examples. It is intended that the following claims be interpreted broadly to embrace all the variations of the preferred embodiments disclosed.